arXiv:1408.2938v1 [cs.CV] 13 Aug 2014

# Learning Multi-Scale Representations for Material Classification

Wenbin Li, Mario Fritz
Max Planck Institute for Informatics, Saarbruecken, Germany
{wenbinli, mfritz}@mpi-inf.mpg.de

## Abstract

*The recent progress in sparse coding and deep learning has made unsupervised feature learning methods a strong competitor to hand-crafted descriptors. In computer vision, success stories of learned features have been predominantly reported for object recognition tasks. In this paper, we investigate if and how feature learning can be used for material recognition. We propose two strategies to incorporate scale information into the learning procedure resulting in a novel multi-scale coding procedure. Our results show that our learned features for material recognition outperform hand-crafted descriptors on the FMD and the KTH-TIPS2 material classification benchmarks.*

## 1. Introduction

Perceiving and recognizing material is a fundamental aspect of visual perception. It enables humans to make predictions about the world and interact with ease. In contrast to texture recognition it requires generalization over large variations between material instances and discriminance between visually similar materials. An efficient material recognition solution will find a wide range of uses such as in context awareness and robot manipulation. Studies have shown that material recognition in real-world scenarios is far from solved [7]. More recently, the task has been pushed even further to less constraint settings. The Flickr Material Dataset (FMD) [20] collects photos from Flickr as samples for common material and demonstrate the difficulties of material recognition. In particular, they incorporate a large number of different descriptors in a Bayesian framework and provides a initial result on the dataset, yet well established manually designed feature descriptor, like LBP [22] and its variants [23, 25] have still been shown to be one of the most powerful methods of feature descriptors and able to achieve state-of-art performance on the material recognition task. It is non-trivial to come up with a good design of visual features and efforts are clearly needed to explore the question how we can automatically learn features for this challenging and relevant problem.

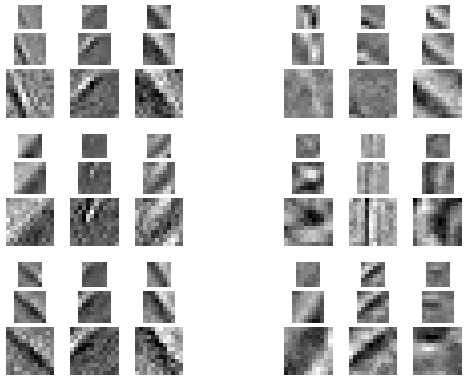

(a) Multi-scale filters learned on the KTH-TIPS2 database.

(b) Multi-scale filters learned on the FMD database.

Figure 1. Example of the multi-scale filters learned on the KTH-TIPS2a (left) and the FMD (right) datasets. These represent a multi-scale code that is learned jointly based on the proposed MS4C model

Recent success of feature learning techniques [13] raises the question if the well established hand-crafted features in material recognition can be replaced with automatically learned ones. It is known that multi-scale representations are key for competitive performance on this task [7, 18, 23]. However, current feature learning techniques do not include multi-scale representations. Therefore, we investigate the applicability of different feature learning techniques to the material recognition task as well as how to bring multi-scale information to the feature learning process.

**Contributions** We present the first study of applying unsupervised feature discovery algorithms for material recognition and show improved performance over hand-crafted feature descriptors. Further, we investigate different ways how to incorporate multi-scale information in the feature learning process. Hereby, we propose the first multi-scale coding procedure that results in a joint representation of multi-scale patches (see Figure 1 for examples of multi-scale codes).

## 2. Related Work

**Material Recognition** Recognition of materials by appearance has received significant attention in the vision community. Curet database [11] was first proposed to address the recognition problem of single material instance, which motivated a lot of progress on texture research [29, 30]. Later research [14, 7] shifted the focus towards whole material class, emphasizing challenges like scale variation and intra-class variation. Liu et al [20] presented the Flickr material dataset which used images from Flickr photos that were captured under unknown real-world conditions. Li et al [19] showed significant improvement over the previous results by only using a simple combination of color and Multi-scale LBP together with rendered data. Hu et al [16] proposed the kernel descriptor [6] and achieved state-of-art performance until recently Qi et al [25] proposed another variant of LBP descriptor to obtain further improvements over previous studies. All these efforts are based on hand designed descriptors while our approach investigates a learning based approach that starts from the raw pixel information.

**Feature Learning** One separate line of research is using learned features to tackle recognition problems. In typical supervised learning setting, one is given a set of examples $X = \{x_1, ..., x_m\}$ and associated labels $\{y_1, ..., y_m\}$. The goal is to learn a model to predict labels for new example $x$. The idea behind unsupervised feature discovery is to find a better representation $\phi(x)$ of data to ease the final learning problem. In the machine learning community, a rich set of models for feature discovery has been proposed. Examples includes sparse coding [26], restricted Boltzmann machines [15, 9] and various autoencoder-based models [4]. The Spike-and-slab sparse coding (S3C) [13] has recently been proposed to combine the advantages of sparse coding with restricted Boltzmann machines and it has shown superior performance. We are based on the S3C model and show how to extend it to multi-scale feature learning as multi-scale feature representation is key in material recognition.

**Multi-Scale Representation** Already, the early texton work included multi-scale filters to enrich the representation. Although the clustering step can be seen as a form of feature learning, the filters are hand-crafted. Also the LBP work has seen extension to a multi-scale LBP [23] that has substantially improved the performance. These feature extraction schemes are entirely hand-crafted. Recently, a multi-scale convolutional neural network (CNN) [12] trained from raw pixels to extract dense feature vectors that encode regions of multiple size centered on each pixel and then performed scene labeling has been proposed. It differs from our multi-scale feature learning approach, as we learn a representation jointly across scales. The image codes derived from our representation directly encode the multi-scaled information. Figure 1 illustrates 12 of such multi-scale codes learned by our model.

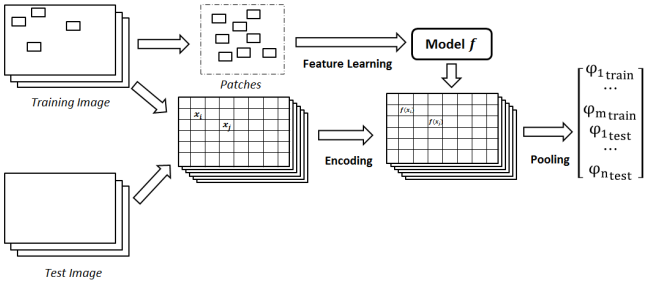


Figure 2. Pipeline for unsupervised feature learning framework.

## 3. Feature Learning

While we have seen broad application and success of feature learning techniques in object recognition, material recognition still relies on hand-crafted features. The appearance of material classes seem special in many ways. First of all, the samples seem to obey a stronger manifold assumption, as the appearance varies rather smoothly w.r.t. changes in lighting direction, orientation and scale. For objects, more drastic changes can occur due to the more pronounced 3D structure and edge information plays an important role.

In this section, we first described the framework for feature learning and then summarize several models we investigated for our tasks. Afterwards, we propose novel multi-scale feature learning strategies in order to accommodate for the multi-scale information that is important for material recognition.

### 3.1. Unsupervised Feature Learning

A commonly used patch-based unsupervised feature learning framework is illustrated in Figure 2. First, random patches $\{v_1, ..., v_n\}$ are extracted from training images and a feature mapping $f$ is learned (dictionary learning). Once the model is obtained, one can encode the patches covering the input image and pool the codes together in order to form the final feature representation (feature extraction). By altering the model used for feature mapping, we can get different feature representations.

**Sparse Coding (SC)** Sparse coding for visual feature coding as illustrated in Figure 3 was originally proposed by Olshauen and Field [24] as an unsupervised learning model of low-level sensory processing in humans. More recently, it was used in the self-taught learning framework [26]. In

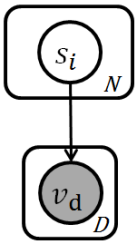


Figure 3. Graphical model describing sparse coding (SC), where $v_d$ denote visible units, $s_i$ represent hidden variables and the directed arrows stand for dependency.

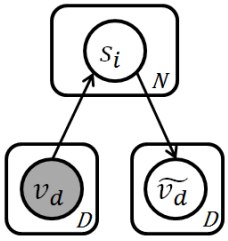


Figure 4. Graphical model describing autoencoder (AE), where $v_d$ denote visible units, $s_i$ for hidden variables, $\widetilde{v_d}$ for reconstructed visible units and the directed arrows stand for dependency.

the first phase, the dictionary $W$ – also known as basis or codes – is obtained by optimizing:

$$\begin{aligned} \underset{W,s_i}{\text{minimize}} \quad & \sum_i \|v_i - W s_i\|_2^2 + \beta \|s_i\|_1 \\ \text{subject to} \quad & \|W_j\|_2 \leq 1, \forall j \end{aligned}$$

Then in the next phase, feature representation $s_i$ for each input $v_i$ is obtained by solving the same form of optimization problem but with the learned dictionary.

**Auto-Encoder (AE)** The Auto-encoder as illustrated in Figure 4 is another popular model widely used for learning feature representation in deep learning community. In the first phase, $v$ is mapped into a latent representation $s$ (encoding) with a nonlinear function $f$ such as the sigmoid function:

$$s = f(Wv + b)$$

Then it is mapped back into a reconstruction $\widetilde{v}$ through a similar transformation:

$$\widetilde{v} = f(\widetilde{W}s + \widetilde{b})$$

and the dictionary (or weights) $W$ is obtained by optimizing the reconstruction error:

$$W = \underset{W}{\operatorname{argmin}} \sum_i L(v_i, \widetilde{v_i})$$

where $L(v_i, \widetilde{v})$ is a loss function such as the squared error $L(v, \widetilde{v}) = \|v - \widetilde{v}\|_2^2$. Then during encoding phase, the features are computed by applying the forward-pass only in oder to obtain $s$.

**Spike-and-Slab Sparse Coding (S3C)** The Spike-and-Slab Sparse Coding (S3C) by Goodfellow at el [13] has been recently proposed to combine the merits of feature learning methods like sparse coding and RBMs.

The model is a two-layer generative process: the first layer is a real-valued $D$-dimensional visible vector $v \in R^D$, where $v_d$ corresponding to the pixel value at position d; the second layer consists of two different kinds of latent variables, the binary *spike* variables $h \in \{0,1\}^N$ and the real-valued *slab* variables $s \in R^N$. The spike variable $h_i$ gates the slab variable $s_i$, and those two jointly define the $i^{th}$ hidden unit as $h_i s_i$. The process can be more formally described as follows:

$$\begin{aligned} \forall i \in \{1,...,N\}, d \in \{1,...,D\} \\ p(h_i = 1) &= g(b_i) \\ p(s_i|h_i) &= N(s_i|h_i\mu_i, \alpha_{ii}^{-1}) \\ p(v_d|s,h) &= N(v_d|W_{d:}(h \circ s), \beta_{dd}^{-1}) \end{aligned}$$

where $g$ is the logistic sigmoid function, $b$ is a set of biases on the spike variables, $\mu$ and $W$ govern the linear dependence of $s$ on $h$ and $v$ on $s$ respectively, $\alpha$ and $\beta$ are diagonal precision matrices of their respective conditionals, and $h \circ s$ denotes the element-wise product of $h$ and $s$. Column of $W$ is constrained to have unit norm, $\alpha$ is restricted to be a diagonal matrix and $\beta$ to be a diagonal matrix or a scalar. In particular, W can be interpreted as a series of filters which can be used sparsely to represent the data. The graphical model describing it is shown in Figure 5 (a).

The model has shown to outperform previous feature learning technique [13] and is the best performer on a recent transfer learning challenge [1] where one trains the model over the patches from the limited number of training images and a large number of unlabeled image data and then coded both training data and test data with the learned model, a standard linear SVM was then used for classification on the learned representation of data.

As discussed in [13], one drawback of sparse coding is that the latent variables are not only encouraged to be sparse, but also to be close to 0 when activated. To tackle

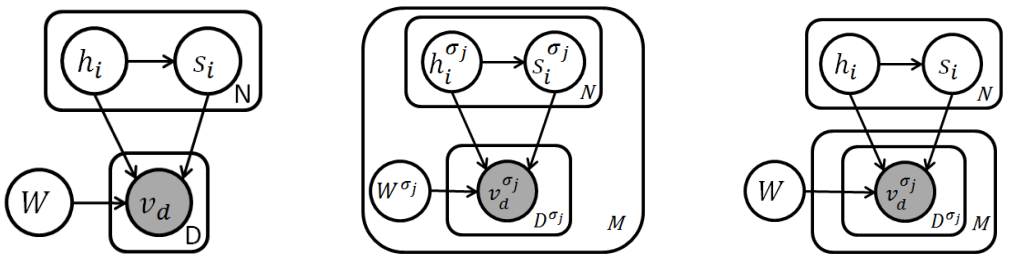


(a) S3C (b) S4C (c) MS4C

Figure 5. Multi-scale feature models. (a) S3C, (b) S4C, (c) MS4C.

this issue, the S3C model introduces separate priors to control the activation of units and the magnitude of activated units separately. Though a similar structured RBM model known as $\mu - ssRBM$ [10] is also proposed for feature learning, the non-factorial posterior of S3C model can grants better discriminative capability by selectively activating only a small set of features for a given input.

**Model Learning** Variational EM [27] algorithm is used for model learning. It is a variant of EM algorithm with modification in the E-step where we only compute a variational approximation to the posterior rather than the posterior itself. In detail, the variational E-step maximize the energy functional with respect to a distribution $Q$ over the unobserved by minimize the Kullback-Leibler divergence: $D_{KL}(Q(h,s)||P(h,s|v)))$, where $Q(h,s)$ is drawn from a restricted family of distributions to ensure that $Q$ is tractable. For more details we refer the reader to [13].

## 4. Multi-Scale Feature Learning

Scale information is a critical element for material and texture recognition problem. [7] showed that explicit treatment of scale is necessary for material recognition in realistic settings. In [19], Li et al performed a manifold alignment with respect to scale between real and synthesized data, which turned to be crucial for using the generated data to improve recognition rate. Similarly, local descriptors like LBP are limited by its small spatial support area, several extensions [23, 21] for multi-scale descriptor have also been shown to yield strong performance improvements. Therefore we propose two different strategies to include multi-scale information in feature learning:

### 4.1. Stacked Spike-and-Slab Sparse Coding (S4C)

In the first strategy, we perform the encoding at multiple scales and stack the obtained codes, then use this code for classification. We convolve the patch with different sized Gaussians before encoding in order to represent scale information. While there is a common dictionary, the representation already encodes how the patch evolves in scale-space and therefore multi-scale information is captured. The graphical model describing it is shown in Figure 5 (b):

$$\begin{aligned}
&\forall i \in \{1,...,N\}, j \in \{1,...,M\}, d \in \{1,...,D\} \\
&p(h_i^{\sigma_j} = 1) = g(b_i^{\sigma_j}) \\
&p(s_i^{\sigma_j}|h_i^{\sigma_j}) = N(s_i^{\sigma_j}|hi^{\sigma_j}\mu_i^{\sigma_j}, (\alpha_{ii}^{\sigma_j})^{-1}) \\
&p(v_d^{\sigma_j}|s^{\sigma_j}, h^{\sigma_j}) = N(v_d^{\sigma_j}|W_{d:}^{\sigma_j}(h^{\sigma_j} \circ s^{\sigma_j}), (\beta_{dd}^{\sigma_j})^{-1})
\end{aligned}$$

where $M$ denotes the number of scales and $\sigma_j$ indexes units and parameters at specific scale.

### 4.2. Multi-Scale Spike-and-Slab Sparse Coding (MS4C)

In the second strategy, we first construct a multi-scale pyramid for each image, apply the feature learning directly on the pyramid and then use the obtained codes for classification. In contrast to the S4C approach, the MS4C approach yields filters/codes that model each patch jointly across scales. . The graphical model describing it is shown in Figure 5 (c):

$$\begin{aligned}
&\forall i \in \{1,...,N\}, j \in \{1,...,M\}, d \in \{1,...,D\} \\
&p(h_i = 1) = g(b_i) \\
&p(s_i|h_i) = N(s_i|hi\mu_i, \alpha_{ii}^{-1}) \\
&p(v_d^{\sigma_j}|s,h) = N(v_d^{\sigma_j}|W_{d:}(h \circ s), \beta_{dd}^{-1})
\end{aligned}$$

where $v_d^{\sigma_j}$ denotes the joint representation of visible units at specific scale $\sigma_j$. Inference is carried out as in the S3C model as the different scales can be seen as a decomposition of a larger multi-scale patch that includes all the scales. Figure 1 shows 12 filters that we have learned in this manner. Each filter reaches across 3 scales.

## 5. Experiments

In our experiments, we investigate how the learning framework can be used for feature discovery on material recognition task and compare our approach to the state-of-the-art on the FMD and the KTH-TIPS2 databases. Further

Figure 6. The KTH-TIPS2a database (left) and the FMD database (right).

we provide insights and visualizations on our learned representations.

### 5.1. Datasets

We use the KTH-TIPS2 database [7] and the Flickr Material Database (FMD) [20] in our experiments. Example images are shown in Figure 6. The KTH-TIPS2 database is designed to study material recognition with a special focus on generalization to novel instance of materials. It includes more than 4000 images from 11 material categories, and each category has 4 different instances. All the instances are imaged from varying viewing angles (frontal, rotated $22.5°$ left and $22.5°$), lighting conditions (from the front, from the side at $45°$, from the top at $45°$, and ambient light) and scales (9 scales equally spaced logarithmically over two octaves), which gives a total of $3 \times 4 \times 9 = 108$ images per instance. We use two instances for training and the other two for test per category. The FMD is collected from Flickr photos, including 10 common material categories with 100 images for each category, 1000 images in total. In our experiment, we randomly split half for training and the other half for testing as suggested in [20].

### 5.2. Experimental setup

We compare the learned features with hand-crafted features by the recognition rates on the two databases with standard SVM classifiers [8]. For single scale experiments, we compare to the LBP [22] and its several variants. For multi-scale approaches we consider: Texton [18], Multi-scale LBP (MLBP) [23]. On the learning side, we compare to vector quantization, sparse coding, auto encoders and the spike-and-slab approach. In particular, we include a comparison with local quantization pattern (LQP)[17]-a recently introduced variant of LBP descriptor and kernel descriptor which has been shown the state-of-art performance on the FMD database. In all our experiments we fix the size of dictionary at 1600 for consistency.

### 5.3. Single-scale

For this group of experiments, we compare the performance between the learned features and the hand-crafted features. In detail, for learned features, we apply the K-means clustering, Auto-Encoder (AE), Sparse Coding(SC) and the S3C model on the patch data where we vary the patch size of $\{6 \times 6, 12 \times 12, 24 \times 24\}$ (we had to skip the results for SC at patch size 24 as it turned out too costly in the encoding phase); for hand-crafted features, we examine the original LBP and several variants of LBP, including uniform-LBP ($LBP_u$), rotation invariant-LBP ($LBP_{ri}$) and rotation invariant, uniform-LBP ($LBP_{ri,u}$) as described in [23].

For implementation, we use Python and base on the library of Theano [5] and Pylearn2 [2] for the auto-encoder and the S3C model, which support GPU computation on network structure. For SC model, we use the SPAMS [3] package.

Experimental results are shown in Figure 7 and Figure 8. Each entry has the results for both the linear kernel (left) and the $exp - \chi^2$ kernel. On both datasets, the S3C model in combination with the linear kernel outperforms all other hand-crafted and learned features. With a performance of $71.3\%$ and $48.4\%$ for the KTH-TIPS2a and the FMD respectively it improves by $4.1\%$ (over $LBP_u$ with the $exp - \chi^2$ kernel) and $9\%$ (over $LBP$ with the linear kernel) respectively. The best performance is achieved for a patch size of 12. We verified that this parameter can be found via cross-validation on the training set. We attribute the decrease in the performance for the patch size of 24 to a lack of data to learn the required number of parameters. Best performance for feature learning technique is typically obtained in combination with linear kernel, while the hand-crafted features have to rely on the non-linear $exp - \chi^2$ kernel. This is another appealing property of the learned features from a computational point of view.

Based on these results, we found that the S3C feature did perform better than other learning approaches and the hand-crafted features for the single-scale setting, and hence we further developed the S3C model to multi-scale approaches in the following experiments.

### 5.4. Multi-scale

For this group of experiment, we introduce scale information with two different models, the Stacked S3C model

(S4C) and the joint Multi-scale S3C model (MS4C), as described in Section 4. In particular, we also investigate the combination of color information for the MS4C model where we concatenate the MS4C codes with the S3C code at the base patch size. For hand-crafted features, we include the Multi-scale LBP (MLBP) and also the texton with the MR8 filter [28]. Though the MR8 filter has been proposed for a long time, it still shows relative good performance on similar recognition tasks [7] and hereby it is also used as a baseline results in our experiments. Furthermore, as the filter banks are manually designed and also contain filters at multiple scales, we count it as a multi-scale hand-crafted feature although the textons are also learned via proper clustering algorithm such as K-means. Experimental results are shown in Figure 7 (a), (c) and Figure 8 (a), (c).

MLBP shows better performance than textons in our experiments. While the S4C model produces slightly worse performance than the MLBP on KTH-TIPS2, we see an improvement of $1.4\%$ for the MS4C. Further including color information improves the performance to $70.5\%$ which is an overall improvement of $3.8\%$ over the best hand-crafted descriptor. From the numbers on the FMD database, we see that our S4C and MS4C beat the best hand-crafted feature (MLBP) by $7.2\%$ and $8\%$ respectively. On this database, inclusion of color information does not yield additional improvements. The new joint multi-scale coding of of the MS4C consistently improves over the stacked approach of the S4C model.

**Further Comparison to State-of-the-Art Descriptors** As not all papers follow the same experimental protocol, we reproduced two additional settings in order to provide more points of comparison to the state-of-the-art. We follow the protocol in [17] and take 3 samples of each class for training and the fourth for testing on the KTH-TIPS2-a data, and then report averages over 4 random partitions via a simple 3-NN classifier, feature learned by single scale S3C at patch size of 12x12 achieved 70.2%, which is significantly better than the reported results of 64.2% for LQP. Also we did additional experiments on the FMD database, following the settings in [16], i.e. performing 5 trials and computing the average, and with multi-scale collaborated representation, we got average recognition rate of 48.3% and standard deviation of 1.8%, which is comparable to the best single kernel descriptor with 49%.

### 5.5. Representation Transfer

Additionally we investigate transferring representations across databases. In detail, we fixed the patch size at 12 and trained single scale S3C on KTH-TIPS2 database and then encoded the FMD data for classification and vice versa. Combined with the results in Figure 7, 8 and Table 1, we can see that when encoding the image data in KTH-TIPS2 with the model learned on FMD, the performance degrades, yet still outperforms single scale LBP and color-patch; when representing the FMD data with the model learned on KTH, the performance even improves over any of the single scale descriptors. This indicates that the features learned through the S3C model on specific dataset are actually eligible to capture some common characteristics which generalize to different data within similar context.

| ClassificationRate(%) | | | | | | |
|---|---|---|---|---|---|---|
| | Single-Scale | | | | Multi-Scale | |
| LBP | $LBP_u$ | $LBP_{ri}$ | $LBP_{ri,u}$ | | Texton | MLBP |
| 58.7/64.8 | 60.3/67.2 | 55.0/53.6 | 50.9/51.4 | | 54.0/58.9 | 66.7/66.1 |

(a) Hand-crafted Feature.

| PatchSize | ClassificationRate(%) | | | |
|---|---|---|---|---|
| | KM | AE | SC | S3C |
| 6 | 60.6/64.8 | 54.3/48.6 | 60.8/64.8 | 63.8/57.5 |
| 12 | 58.4/65.5 | 49.6/44.2 | 66.0/64.8 | 71.3/66.0 |
| 24 | 58.3/65.0 | 48.9/39.1 | * | 55.9/60.8 |

(b) Standard Feature Learning.

| ClassificationRate(%) | | |
|---|---|---|
| S4C | MS4C | MS4C+Color |
| 65.6/58.6 | 68.1/66.6 | 70.5/69.3 |

(c) Multi-scale Feature Learning.

Figure 7. Experimental results on the KTH-TIPS2a database, where the first number in each entry is the recognition rate with the linear kernel, the other one is the number with the $exp - \chi^2$ kernel.

| ClassificationRate(%) | | | | | |
|---|---|---|---|---|---|
| Single-Scale | | | | Multi-Scale | |
| LBP | $LBP_u$ | $LBP_{ri}$ | $LBP_{ri,u}$ | Texton | MLBP |
| 39.4/36 | 38.2/36.2 | 34.2/35.6 | 27.8/31.8 | 29.4/35.6 | 41.4/42.0 |

(a) Hand-crafted Feature.

| PatchSize | ClassificationRate(%) | | | |
|---|---|---|---|---|
| | KM | AE | SC | S3C |
| 6 | 29.2/38.0 | 37.6/25.0 | 34.8/30.8 | 42.6/39.2 |
| 12 | 26.0/39.6 | 32.4/25.0 | 39.4/26.4 | 48.4/41.8 |
| 24 | 26.8/37.2 | 29.2/22.0 | * | 40.8/44.0 |

(b) Standard Feature Learning.

| ClassificationRate(%) | | |
|---|---|---|
| S4C | MS4C | MS4C+Color |
| 49.2/42.2 | 50.0/41.0 | 48.8/43.2 |

(c) Multi-scale Feature Learning.

Figure 8. Experimental results on FMD database, where the first number in each entry is the recognition rate with the linear kernel, the other one is the number with the $exp - \chi^2$ kernel.

| Feature | ClassificationRate(%) | |
|---|---|---|
| | Linear Kernel | $exp - \chi^2$ Kernel |
| Code learned on FMD, and represent KTH-TIPS2 | | |
| S3C($12 \times 12$) | 65.98 | 61.11 |
| Code learned on KTH-TIPS2, and represent FMD | | |
| S3C($12 \times 12$) | 44.4 | 43.2 |

Table 1. Results for Transfer Representation.

| | | ClassificationRate(%) | |
|---|---|---|---|
| | | S3C | MS4C |
| Scale | {5} | 64.7/64.3 | 65.2/58.4 |
| | {3,7} | 65.8/66.4 | 67.7/65.9 |
| | {3,5,7} | 65.1/66.7 | 70.7/67.7 |
| | {1,3,5,7,9} | 67.3/66.4 | 67.7/66.5 |
| | {1,2,3,4,5,6,7,8,9} | 71.3/66.0 | 68.1/66.6 |

Table 2. Recognition rates with different number scales available in the training dataset on KTH-TIPS2a database. The original database covers 9 different scales spanning two octaves which are indexed by No. 1 to 9 here.

### 5.6. Discussion

**Visualization of Models** Figure 1 shows visualization of our proposed Multi-Scale Spike-and-Slab Sparse Coding model. We see how each filter has a multi-scale response. We looked at a larger range of such filters, which reveals some more interesting properties. Some of these filters have a very similar structure across scales, while other do vary strongly. This observation and the strong performance numbers in our experiments let us conclude that a multi-scale code indeed captures additional information about how edge structures propagate through scales.

**Effect of Patch Size** Feature learning results for single scale descriptors are dependent on the patch size. The size of patch determines the locality of the descriptor and therefore affects how the descriptor can be generalized to different instances, and from our experience, there seems not be any overall optimal patch size, which suggests we may need to try several candidates for a specific dataset and select the best one for use. In our experiments, we found that the patch size can be chosen based on cross-validation on the training set. Furthermore, our multi-scale approach S4C and MS4C resolve this problem by learning the representation across multi-scales.

**Scale Information** Most of time, we see improvements when incorporating scale information, however on the KTH-TIPS2 database, we find that a descriptor learned at single scale performs the best. This may be related to the properties of the specific dataset and also the nature of our designed multi-scale descriptor. Both strategies for our multi-scale descriptors involve some redundancy between every scale that may degrade the classification performance, in return, this redundancy also encodes the scale information by itself that could improve performance, and the final performance will be affected by these two factors jointly. As for the KTH-TIPS2 database, material images were taken under strictly controlled conditions, in particular, only 9 different scales for all the instances, so the improvement via scale information is very limited in this case while the redundancy still affect the classification rate negatively. This in particular explains why the two multi-scale descriptors which already incorporate the information in model with patch size of 12 get worse results than the single-scale descriptor. In contrast, for the case of the FMD database, images were collected from Flickr photos in arbitrary conditions, scale information become significantly more important and surpass the influence from the redundancy, which makes the multi-scale descriptor beat any of its components at single scale. In real world application, it seems closer to the latter situation, thus the multi-scale descriptor is preferable in this sense.

To further validate our analysis, we design additional experiments by making use of only subset of training data on the KTH-TIPS2 data in order to align the settings with the FMD database. So compared with the standard settings in the KTH-TIPS2 database where training data covering all different scales appeared in both training and test partition, we only use images taken under some scales for feature learning. Note this setting also resembles the situation in the FMD database, where data are imaged under unknown conditions such that training data cannot include the same scale information in test data. In this way, we would like to see if the multi-scale feature learning can provide extra power over the basic feature learning model. The experimental results are shown in Table 2. As we can see from the table, when observed very limited scales in training data like only one to three scales, the MS4C indeed outperforms the basic S3C single scale model.

**Color vs. Gray-scale** While color information serves as an important cue for visual recognition, it could also lead to confusion, so we should be careful to incorporate the color information. It is interesting to compare the results for the two multi-scale joint representation, with one in gray-scale and the other in color: on the KTH-TIPS2 database, the color information led to an improvement over the gray-scale representation while the gray-scale version achieved the best performance on the FMD database. It could be explained by the large variation of color information in the FMD data which causes the confusion, whereas the color cue is simpler and more informative for classification on KTH-TIPS2.

## 6. Conclusions

We have investigated different feature learning strategies for the task of material classification. Our results match and even surpass standard hand-crafted descriptors. Furthermore, we extended feature learning techniques to incorporate scale information. We propose the first coding procedure that learns and encodes features with a joint multi-scale representation. The comparison of our learned features with state-of-the-art descriptors shows improved performance on standard material recognition benchmarks.

## References


[1] Challenges in learning hierarchical models: Transfer learning and optimization. https://sites.google.com/site/nips2011workshop/transfer-learning-challenge. 3

[2] Pylearn2 vision, a python library for machine learning. http://deeplearning.net/software/pylearn2/. 5

[3] Sparse modeling software, an optimization toolbox for solving various sparse estimation problems. http://spams-devel.gforge.inria.fr/. 5

[4] Y. Bengio, P. Lamblin, D. Popovici, and H. Larochelle. Greedy layer-wise training of deep networks. In *NIPS*, 2007. 2

[5] J. Bergstra, O. Breuleux, F. Bastien, P. Lamblin, R. Pascanu, G. Desjardins, J. Turian, D. Warde-Farley, and Y. Bengio. Theano: a CPU and GPU math expression compiler. In *Proceedings of the Python for Scientific Computing Conference (SciPy)*, 2010. 5

[6] L. Bo, X. Ren, and D. Fox. Kernel descriptors for visual recognition. In *NIPS*, 2010. 2

[7] B. Caputo, E. Hayman, and P. Mallikarjuna. Class-specific material categorisation. In *ICCV*, 2005. 1, 2, 4, 5, 6

[8] C.-C. Chang and C.-J. Lin. LIBSVM: A library for support vector machines. *ACM Transactions on Intelligent Systems and Technology*, 2011. 5

[9] A. Courville, J. Bergstra, and Y. Bengio. A spike and slab restricted boltzmann machine. *JMLR*, 2011. 2

[10] A. Courville, J. Bergstra, and Y. Bengio. Unsupervised models of images by spike-and-slab rbms. In *ICML*, 2011. 3

[11] K. J. Dana, B. van Ginneken, S. K. Nayar, and J. J. Koenderink. Reflectance and texture of real-world surfaces. *ACM Trans. Graph.*, 1999. 2

[12] C. Farabet, C. Couprie, L. Najman, and Y. LeCun. Scene Parsing with Multiscale Feature Learning, Purity Trees, and Optimal Covers. In *ICML*, 2012. 2

[13] I. Goodfellow, A. Couville, and Y. Bengio. Large-scale feature learning with spike-and-slab sparse coding. In *ICML*, 2012. 1, 2, 3, 4

[14] E. Hayman, B. Caputo, M. Fritz, and J.-O. Eklundh. On the significance of real-world conditions for material classification. In *ECCV*, 2004. 2

[15] G. Hinton, S. Osindero, and Y. Teh. A fast learning algorithm for deep belief nets. *Neural computation*, 2006. 2

[16] D. Hu, L. Bo, and X. Ren. Toward robust material recognition for everyday objects. In *BMVC*, 2011. 2, 6

[17] S. U. Hussain and B. Triggs. Visual recognition using local quantized patterns. In *ECCV*, 2012. 5, 6

[18] T. Leung and J. Malik. Representing and recognizing the visual appearance of materials using three-dimensional textons. *IJCV*, 2001. 1, 5

[19] W. Li and M. Fritz. Recognizing materials from virtual examples. In *ECCV*, 2012. 2, 4

[20] C. Liu, L. Sharan, E. H. Adelson, and R. Rosenholtz. Exploring features in a bayesian framework for material recognition. In *CVPR*, 2010. 1, 2, 5

[21] T. Mäenpää and M. Pietikäinen. Multi-scale binary patterns for texture analysis. *Image Analysis*, 2003. 4

[22] T. Ojala, M. Pietikäinen, and D. Harwood. A comparative study of texture measures with classification based on featured distributions. *Pattern Recognition*, 1996. 1, 5

[23] T. Ojala, M. Pietikainen, and T. Maenpaa. Multiresolution gray-scale and rotation invariant texture classification with local binary patterns. *TPAMI*, 2002. 1, 2, 4, 5

[24] B. A. Olshausen et al. Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature*, 1996. 2

[25] X. Qi, R. Xiao, J. Guo, and L. Zhang. Pairwise rotation invariant co-occurrence local binary pattern. In *ECCV*, 2012. 1, 2

[26] R. Raina, A. Battle, H. Lee, B. Packer, and A. Ng. Self-taught learning: transfer learning from unlabeled data. In *ICML*, 2007. 2

[27] L. Saul and M. Jordan. Exploiting tractable substructures in intractable networks. 1996. 4

[28] M. Varma and A. Zisserman. Classifying images of materials: Achieving viewpoint and illumination independence. In *Computer VisionłECCV 2002*. 2002. 6

[29] M. Varma and A. Zisserman. A statistical approach to texture classification from single images. *IJCV*, 2005. 2

[30] M. Varma and A. Zisserman. A statistical approach to material classification using image patch exemplars. *TPAMI*, 2009. 2